arXiv:1607.07343v1 [math.ST] 25 Jul 2016

Gaussian processes and Bayesian moment estimation*

Jean-Pierre Florens† Anna Simoni‡

Toulouse School of Economics CREST and Ensae

This version: July 2016

Abstract

Given a set of estimating equations (EE) that characterize a parameter θ, we investigate a semiparametric Bayesian approach for inference on θ that does not restrict the data distribution F apart from the EE. As main contribution, we construct a degenerate Gaussian process prior that, conditionally on θ, restricts the F generated by this prior to satisfy the EE with probability one. Our prior works even in the more involved case where the number of EE is larger than the dimension of θ. We show that this prior is computationally convenient. Since the likelihood function is not specified by the model, we approximate it based on a linear functional transformation of F that has an asymptotically Gaussian empirical counterpart. This likelihood is used to construct the posterior distribution. We provide a frequentist validation of our procedure by showing consistency and asymptotic normality of the posterior distribution of θ.

Key words: Estimating Equations, Gaussian processes, overidentification, invariance, posterior consistency, functional equation.

*First version: February 2012. The authors especially thank Yuichi Kitamura for insightful discussions. The authors are grateful to seminars and conferences in: Berlin, Boston College, Bristol, Carlos III, CREST, Northwestern, SBIES 2015, ICEEE 2015, NASM 2012, Toulouse. We thank financial support from ANR-13-BSH1-0004 (IPANEMA). Anna Simoni gratefully acknowledges financial support from Labex ECODEC (ANR - 11-LABEX-0047), SFB-884 and hospitality from the University of Mannheim.

†Toulouse School of Economics - 21, allée de Brienne - 31000 Toulouse (France). Email: jean-pierre.florens@tse-fr.eu

‡CREST, 15 Boulevard Gabriel Péri, 92240 Malakoff (France). Email: simoni.anna@gmail.com (*corresponding author*).

1 Introduction

Statistical models are often formulated via estimating equations (EE) - or moment restrictions - of the form $\mathbf{E}^F[h(\theta, X)] = 0$, where $h(\theta, X)$ is a vector-valued function of an observable random element X with distribution F and a parameter vector θ. These EE provide the only information available about θ and the data distribution. Given a set of EE, this paper builds a semiparametric Bayesian inference procedure for θ that imposes these moment conditions in the nonparametric prior distribution for the data distribution F and that is computationally convenient. Apart from these moment restrictions, F is left unrestricted.

A main advantage of Bayesian inference consists in providing a well-defined posterior distribution that is important for many decision problems and for predictive analysis. On the other hand, constructing Bayesian inference procedures for moment condition models presents two difficulties. A first difficulty is due to the fact that a likelihood is not available. A second difficulty arises when the number of EE exceeds the dimension of θ (overidentification) because imposing overidentifying moment restrictions on the prior distribution for the nonparametric F is challenging. The contribution of this paper is to propose an elegant approach that allows to deal with these two difficulties. As a by-product we show that the quasi-likelihood of some Laplace-type procedures arises as the limit of our Bayesian procedure when our prior becomes diffuse.

The model we consider is as follows. Let X be an observable random element in $\mathbb{R}^m$ with distribution F and $X_1, \ldots, X_n$ be an *i.i.d.* sample of X. The parameter $\theta \in \Theta \subset \mathbb{R}^p$ is linked to the data generating process (DGP) F through the EE

$$\mathbf{E}^F [h(\theta, X)] = 0, \tag{1.1}$$

where $h(\theta, x) := (h_1(\theta, x), \ldots, h_d(\theta, x))^T$ and the functions $h_j(\theta, x)$, $j = 1, \ldots, d$ are real-valued and known up to θ. We assume $d \geq p$ and our main interest is the case where $d > p$, which is in general more challenging than the case $d = p$. Examples of the case $d > p$ can be found in Broniatowski and Keziou [2012]. Apart from (1.1), F is completely unrestricted. The Bayesian procedure proposed in this paper constructs a nonparametric prior distribution for F with support equal to the subset of distributions that satisfy the EE for a given θ.

Imposing moment restrictions via semiparametric priors may be challenging depending on the relationship existing between θ and F. More precisely, when the model is exactly identified (*i.e.* $p = d$) and (1.1) characterizes θ as an explicit function of F, say $\theta = b(F)$ for some function $b(\cdot)$, then (1.1) does not restrict F. Therefore, one places an unrestricted nonparametric prior on F and recovers the prior of θ via the transformation $\theta = b(F)$. The (θ, F)s generated by this prior automatically satisfy the constraints.

On the contrary, when $d > p$ (overidentified model), θ cannot be expressed as an explicit function of F. Indeed, (1.1) imposes constraints on F and the existence of a solution θ to (1.1) is guaranteed only for a subset of distributions F. Therefore, a restricted nonparametric prior on F must be specified conditionally on θ and the support of this prior is a proper subset of the set of probability distributions. It turns out that incorporating overidentifying moment restrictions in a semiparametric prior for (θ, F) is not straightforward. In this paper we propose a way to construct a semiparametric prior that incorporates the overidentifying EE.

Our strategy is based on a degenerate Gaussian process ($\mathcal{GP}$) prior with restricted support which is easy to deal with and that works as follows. The DGP F is assumed to admit a density function f with respect to some positive measure Π chosen by the researcher (for instance the Lebesgue measure). Then, we endow f with a $\mathcal{GP}$ prior conditional on θ. The $d > p$ EE are incorporated by constraining the prior mean and prior covariance of this $\mathcal{GP}$ in an appropriate way. Because this prior imposes the moment restrictions, it is degenerate on a proper subset of the set of probability density functions. The reason for the appropriateness of a $\mathcal{GP}$ prior in such a framework is due to the fact that the EE in (1.1) are linear in f and the linearity of the model matches extremely well with a $\mathcal{GP}$ prior. A remarkable feature of our method is that the EE are imposed directly through the $\mathcal{GP}$ prior of f given θ without requiring a second step projection over the set of density functions satisfying the moment restrictions. To the best of our knowledge a $\mathcal{GP}$ prior has not been used yet in the EE framework.

Our Bayesian procedure, that we call the $\mathcal{GP}$-approach, is constructed as follows. In the overidentified case ($d > p$) we first specify a prior on θ and then a $\mathcal{GP}$ prior on f conditional on θ. In the case $d = p$ we may either proceed as in the overidentified case or specify an unrestricted $\mathcal{GP}$ prior on f and then deduce from it the prior for θ through the explicit relationship $\theta = b(f)$. We circumvent the difficulty of the likelihood function specification, which is not available, by constructing a linear functional transformation of the DGP F such that its empirical counterpart, say r_n, has an asymptotic Gaussian distribution. This will be used as the sampling model. Therefore, our model is approximately conjugate and allows easy computations while being nonparametric in F.

We provide a closed-form expression for the marginal posterior distribution of θ (obtained by integrating out f) and propose the maximum of this distribution as an estimator for θ. The maximum a posteriori of θ is usually not available in closed-form but can be easily computed via drawn from the marginal posterior. We show that the quasi-likelihood function – also called limited information likelihood (LIL) – used, among others, by Kim [2002] and Chernozhukov and Hong [2003], can be obtained as the limit of our marginal likelihood when the $\mathcal{GP}$ prior for f is allowed to become diffuse. In addition, when the prior for f becomes noninformative, the marginal posterior distribution for θ becomes the same

(up to constants) as the GEL objective function with quadratic criterion and is a monotonic transformation of the continuous updating GMM objective function (Hansen et al. [1996]).

Finally, we provide a frequentist validation of our method by showing: *(i)* posterior consistency, *(ii)* frequentist consistency of the maximum a posteriori estimator, and *(iii)* asymptotic normality of the posterior distribution of θ.

Related literature. Inference in an EE framework has received very much attention in the past literature. Among the most popular statistical methods in this framework is the Empirical Likelihood (EL)-based inference (*e.g.* Owen [1988], Qin and Lawless [1994], Imbens, Owen [2001], Chen and Van Keilegom [2009]) and its variations like the Exponential Tilting (*e.g.* Kitamura and Stutzer [1997] and Kitamura [1997]), the Exponentially Tilted EL (ETEL) proposed by Schennach [2007], and the Generalized EL (GEL) (*e.g.* Smith [1997], Newey and Smith [2004]). EL and ETEL can also be used in a Bayesian framework for valid inference as established by Lazar [2003], Grendar and Judge [2009], Schennach [2005], Rao and Wu [2010], Chib et al. [2016] and Chaudhuri et al. [2017].

An alternative frequentist inference procedure, which is very popular especially in econometrics, is the Generalized method of Moments (GMM) (Hansen [1982] and Hansen and Singleton [1982]). From a Bayesian inference point of view, alternative procedures that have been proposed in the literature can be classified in two types depending on whether the moment conditions (1.1) are imposed in the likelihood or in the nonparametric prior. Procedures of the first type construct a quasi-likelihood by exponentiating the quadratic criterion (or the GEL criterion) associated with the empirical counterpart of (1.1) and include Kwan [1999], Kim [2002], Chernozhukov and Hong [2003] and Liao and Jiang [2011] among others. Our paper shows that the quasi-likelihood used in this type of approach arises as the limit of our $\mathcal{GP}$ prior as it becomes diffuse. We provide thus a fully Bayesian justification to this approach. Procedures of the second type impose the moment conditions in the prior for (θ, F) while leaving the likelihood completely unrestricted and include, in addition to the Bayesian ET and ETEL discussed above, Chamberlain and Imbens [2003] who use a Dirichlet prior, Kitamura and Otsu [2011] and Shin [2014] who propose a two-step procedure based on a projection of a Dirichlet process mixture prior and of a mixture of Dirichlet Process prior, respectively, and Bornn et al. [2015] who use Hausdorff measures to build probability tools for dealing with moment estimation.

The paper is organized as follows. The $\mathcal{GP}$-approach is described in section 2, which contains our main contribution. Here, we also show the link existing between our approach and some frequentist approaches in an EE framework. In section 3 we analyze asymptotic properties of the posterior distribution of θ and of the maximum a posteriori estimator. In section 4 we detail how to implement our method through simulation studies. All the proofs are gathered in the Appendix which is contained in the Supplementary Material.

2 The Gaussian Process ($\mathcal{GP}$)-approach

Let X be a continuous random element in $S \subseteq \mathbb{R}^m$ with distribution F and $X_1, \ldots, X_n$ be an i.i.d. sample of X. Assume that F is absolutely continuous with respect to some positive measure Π (*e.g.* the Lebesgue measure) with density function f. In other words, conditionally on f the data are drawn from F: $X_1, \ldots, X_n | f \sim F$. The set of probability density functions on S with respect to Π is denoted by M.

Let $\theta \in \Theta \subseteq \mathbb{R}^p$ be the parameter of interest characterized by (1.1). By adopting a frequentist point of view, we denote, throughout the paper, the true value of θ by θ_*, the true DGP by F_* and its density with respect to Π by f_*. The model is assumed to be well-specified, that is, there exists $\theta_* \in \Theta$ such that $\mathbf{E}^{F_*}(h(\theta_*, X)) = 0$ holds. We endow $S \subseteq \mathbb{R}^m$ with the trace of the Borelian σ-field $\mathfrak{B}_S$ and specify Π as a positive measure on this subset. We denote by $\mathcal{E} = L^2(S, \mathfrak{B}_S, \Pi)$ the Hilbert space of square integrable functions on S with respect to Π and by $\mathfrak{B}_\mathcal{E}$ the Borel σ-field generated by the open sets of $\mathcal{E}$. The scalar product and norm on this space are defined in the usual way and denoted by $\langle \cdot, \cdot \rangle$ and $|| \cdot ||$, respectively.

The parameters of the model are (θ, f), where f is the nuisance parameter, and the parameter space is

$$\Lambda = \left\{ (\theta, f) \in \Theta \times \mathcal{E}_M; \int h(\theta, x) f(x) \Pi(dx) = 0 \right\}, \qquad \mathcal{E}_M := \mathcal{E} \cap M,$$

where $h : \Theta \times \mathbb{R}^m \to \mathbb{R}^d$ is a known function. In the following of the paper we maintain the following assumption.

Assumption 2.1. *(i) The true f_* satisfies $f_* \in \mathcal{E}_M := \mathcal{E} \cap M$; (ii) the moment function $h(\theta, \cdot)$ is such that $h_i(\theta, \cdot) \in \mathcal{E}$ for every $i = 1, \ldots, d$ and for every $\theta \in \Theta$, where h_i denotes the i-th component of h; (iii) $d \geq p$.*

Assumption 2.1 *(i)* restricts f_* to be square integrable with respect to Π and is for instance verified if f_* is bounded and Π is a bounded measure. The model is made up of three elements that we detail in the next two subsections: a prior on θ, denoted by $\mu(\theta)$, a conditional prior on f given θ, denoted by $\mu(f|\theta)$, and the sampling model. In the following, we shorten "almost surely" by "a.s." and omit the probability which "a.s." refers to. We denote by $\mathbf{E}^F$ the expectation taken with respect to F and by $\mathbf{E}^*$ the expectation taken with respect to F_*.

2.1 Prior distribution

We specify a prior probability measure μ for (θ, f) of the form $\mu(\theta, f) = \mu(\theta)\mu(f|\theta)$. By abuse of notation, $\mu(\theta)$ will also denote the Lebesgue density of the prior distribution of θ

in the case it admits it. The prior $\mu(\theta)$ may either be flat (non-informative) or incorporate any additional information available to the econometrician about θ. In any case, it is tacitly assumed that $\mu(\theta)$ is such that the posterior of θ exists.

Given a value for θ, the conditional prior $\mu(f|\theta)$ is specified such that its support equals the subset of functions in $\mathcal{E}$ that integrate to one and satisfy (1.1) for this particular value of θ. At the best of our knowledge, the construction of such a conditional prior $\mu(f|\theta)$ is new in the literature and we now explain it in detail.

Construction of the conditional prior $\mu(f|\theta)$. We construct the conditional prior distribution $\mu(f|\theta)$ of f, given θ, as a $\mathcal{GP}$ on $\mathfrak{B}_{\mathcal{E}}$ with mean function $f_{0\theta} \in \mathcal{E}_M$ and covariance operator $\Omega_{0\theta} : \mathcal{E} \to \mathcal{E}$. We restrict $f_{0\theta}$ and $\Omega_{0\theta}$ to guarantee that the trajectories f generated by $\mu(f|\theta)$ are such that the corresponding F (which is such that $dF = f d\Pi$) integrates to 1 and satisfies equation (1.1) with probability 1. The two sets of restrictions that we impose are the following (one on $f_{0\theta}$ and one on $\Omega_{0\theta}$):

Restriction 1 (Restriction on $f_{0\theta}$)**.** The prior mean function $f_{0\theta} \in \mathcal{E}_M$ is chosen such that

$$\int h(\theta, x) f_{0\theta}(x) \Pi(dx) = 0. \tag{2.1}$$

Restriction 2 (Restriction on $\Omega_{0\theta}$)**.** The prior covariance operator $\Omega_{0\theta} : \mathcal{E} \to \mathcal{E}$ is chosen such that

$$\begin{cases} \Omega_{0\theta}^{1/2} h(\theta, x) &= 0 \\ \Omega_{0\theta}^{1/2} 1 &= 0 \end{cases} \tag{2.2}$$

where $\Omega_{0\theta}^{1/2} : \mathcal{E} \to \mathcal{E}$ denotes the positive square root of $\Omega_{0\theta}$: $\Omega_{0\theta} = \Omega_{0\theta}^{1/2} \Omega_{0\theta}^{1/2}$.

The covariance operator $\Omega_{0\theta}$ is linear, self-adjoint and trace-class.[1] Due to Restriction 2, $\Omega_{0\theta}$ is not injective. In fact, the null space of $\Omega_{0\theta}$, denoted by $\mathfrak{N}(\Omega_{0\theta})$, is not trivial and contains effectively the constant 1 – which implies that the trajectories f generated by the prior integrate to 1 a.s. (with respect to Π) – and the functions $h(\theta, x)$ – which implies that the trajectories f satisfy the moment conditions a.s. This means that $\Omega_{0\theta}$ is degenerate in the directions along which we want that the corresponding projections of f and $f_{0\theta}$ are equal. Therefore, the support of $\mu(f|\theta)$ is a proper subset of $\mathcal{E}$. This is the meaning of the next lemma.

Lemma 2.1. *The conditional $\mathcal{GP}$ prior $\mu(f|\theta)$, with mean function $f_{0\theta}$ and covariance operator $\Omega_{0\theta}$ satisfying Restrictions 1 and 2, generates trajectories f that satisfy $\mu(f|\theta)$-a.s. the conditions*

[1] A trace-class operator is a compact operator with eigenvalues that are summable. Remark that this guarantees that the trajectories f generated by $\mu(f|\theta)$ satisfy $\int f^2 d\Pi < \infty$ a.s.

$$\int f(x)\Pi(dx) = 1 \quad \textit{and} \quad \int h(\theta, x) f(x)\Pi(dx) = 0.$$

Remark 2.1. Restrictions 1 and 2 imply that the trajectories generated by $\mu(f|\theta)$ integrates to 1 (with respect to Π) and satisfy (1.1) a.s. but they do not guarantee non-negativity of the trajectories. Thus, the support of $\mu(f|\theta)$ is smaller than $\mathcal{E}$ but bigger than $\mathcal{E}_M$. To impose non-negativity one could: (i) either project the prior on the space of non-negative functions, or (ii) write $f = g^2$ or $f = e^g / \int e^g$, $g \in \mathcal{E}$, and specify a conditional prior distribution for g, given θ, instead of for f. Nonetheless, it is important to notice that: the projected prior in (i) is not Gaussian anymore, and in (ii) we cannot use our restricted $\mathcal{GP}$ prior because if it is specified on g it does not work to impose the moment restrictions on f so that a different restricted prior should be constructed. Moreover, in both (i) and (ii) the resulting posterior for θ is not available in closed form which is instead one of the main advantages of our procedure. Therefore, it is not possible to impose the non-negativity constraint if one wants to use our restricted $\mathcal{GP}$ prior. However, because our goal is to make inference on θ while f is a nuisance parameter, failing to impose the non-negativity constraint is not an issue as long as our procedure is shown to be consistent for θ (which we show in section 3).

From a practical implementation point of view, a covariance operator satisfying Restriction 2 and a $f_{0\theta}$ satisfying Restriction 1 may be constructed as follows.

Construction of $\Omega_{0\theta}$. Let $(\lambda_j)_{j\in\mathbb{N}}$ be a decreasing sequence of non-negative numbers such that $\sum_j \lambda_j < \infty$, and $(\varphi_j)_{j\in\mathbb{N}}$ be an orthonormal basis (o.n.b.) for $\mathcal{E}$. Then, $\forall \phi \in \mathcal{E}$: $\Omega_{0\theta}\phi = \sum_{j=0}^{\infty} \lambda_j \langle \phi, \varphi_j \rangle \varphi_j$. Remark that $(\lambda_j)_{j\in\mathbb{N}}$ and $(\varphi_j)_{j\in\mathbb{N}}$ correspond to the eigenvalues and eigenfunctions of $\Omega_{0\theta}$, respectively. Since the null space $\mathfrak{N}(\Omega_{0\theta}) \subset \mathcal{E}$ is spanned by $\{1, h_1(\theta,\cdot), \ldots, h_d(\theta,\cdot)\}$, we can set the first eigenfunctions of $\Omega_{0\theta}$ equal to the elements of any o.n.b. of $\mathfrak{N}(\Omega_{0\theta})$. Restriction 2 is then fulfilled by setting the corresponding eigenvalues equal to 0. For instance, if $\{1, h_1(\theta,\cdot), \ldots, h_d(\theta,\cdot)\}$ are orthonormal as elements of $\mathcal{E}$, then $\mathfrak{N}(\Omega_{0\theta})$ has dimension $d+1$, the first eigenfunctions are $(\varphi_0, \varphi_1, \ldots, \varphi_d)^T = (1, h^T)^T$ and the corresponding eigenvalues are $\lambda_j = 0$, $\forall j = 0, 1, \ldots, d$. Remark that in this case, necessarily, $\int \Pi(dx) = 1$. If $\{1, h_1(\theta,\cdot), \ldots, h_d(\theta,\cdot)\}$ are not orthonormal then one can use their orthonormalized counterparts as the first eigenfunctions of $\Omega_{0\theta}$. The latter is the method we use to implement our procedure. The remaining components $(\varphi_j)_{j>d}$ are chosen such that $(\varphi_j)_{j\geq 0}$ forms an orthonormal basis of $\mathcal{E}$ and $(\lambda_j)_{j>d}$ are chosen such that $\sum_{j>d} \lambda_j < \infty$. Hence,

$$\forall \phi \in \mathcal{E}, \quad \Omega_{0\theta}\phi = \sum_{j=d+1}^{\infty} \lambda_j \langle \phi, \varphi_j \rangle \varphi_j$$

where we suppress the dependence of φ_j on θ for simplicity. Examples of choices for $(\lambda_j)_{j>d}$ are, for some constant $c > 0$: *(i)* $\lambda_j = cj^{-a}$ with $a > 1$, *(ii)* $\lambda_j = ce^{-j}$. In section 4 we

clarify the construction of $\Omega_{0\theta}$.

Construction of $f_{0\theta}$. By definition, $f_{0\theta}(\cdot) \in \mathcal{E}$ is orthogonal to $h(\theta, \cdot)$. Therefore, if the functions $\{\varphi_j\}_{j>d}$ used to construct $\Omega_{0\theta}$ are bounded in absolute value on S with bounds $\{B_j\}_{j>d}$, respectively, then we can select a sequence of constants $\{a_j\}_{j>d}$ that satisfy $\sum_{j>d} |a_j| B_j \leq 1$ and construct $f_{0\theta}$ as $f_{0\theta}(x) = 1 + \sum_{j>d} a_j \varphi_j(x)$. Alternatively, if not all the $\{\varphi_j\}_{j>d}$ are bounded we can set all but a finite number of the $\{a_j\}_{j>d}$ to zero.

Remark 2.2. In the just-identified case where $d = p$ and the EE (1.1) can be solved explicitly for θ (that is, $\theta = b(f)$, for some functional b), then the prior for (θ, f) may be constructed in an alternative way: one can first specify an unrestricted prior for f and then recover from it the prior for θ. When b is a linear functional and θ can take any value in $\mathbb{R}^p$, one can specify a $\mathcal{GP}$ prior $\mu(f)$ for f (independent of θ) with a mean function f_0 restricted only to be a *pdf* and a covariance operator Ω_0 restricted only to satisfy $\Omega_0^{1/2} 1 = 0$. Then, the prior for θ is obtained through the transformation $b(\cdot)$ and will be Gaussian. Because the support of this prior is $\mathbb{R}^p$, then this approach is feasible if every value in $\mathbb{R}^p$ is plausible for θ. For example, if $\theta = \mathbf{E}^F(x)$ and the support of x is $\mathbb{R}^p$, then $b(f) = \langle f, \iota \rangle$ and $\mu(\theta) = \mathcal{N}(\langle f_0, \iota \rangle, \langle \Omega_0 \iota, \iota \rangle)$, where $\iota \in \mathcal{E}$ denotes the identity function $\iota(x) = x$.

2.2 The sampling model

Given the observed *i.i.d.* sample $(x_1, \ldots, x_n)$, the likelihood function is $\prod_{i=1}^n f(x_i)$. While apparently simple, using this likelihood for Bayesian inference on θ makes the analysis of the posterior distribution complicated. This is because to compute the posterior for θ one has to marginalize out f. Since a $\mathcal{GP}$ prior is not a natural conjugate of the *i.i.d.* model then, marginalization of f has to be carried out through numerical, or Monte Carlo, integration on a functional space, which may be computationally costly. To avoid this difficulty, we propose an alternative and original way to construct the sampling model that allows for a conjugate analysis and prevents from numerical integration. Our approach is based on a functional transformation r_n of the sample $X_1, \ldots, X_n$.

This transformation r_n is chosen by the researcher and must have the following characteristics: *(I)* r_n is an observable element of an infinite-dimensional Hilbert space $\mathcal{F}$ (to be defined below), for instance a L^2-space; *(II)* r_n converges weakly towards a Gaussian process in $\mathcal{F}$; *(III)* the expectation of r_n, conditional on f, defines a linear operator $K : \mathcal{E} \to \mathcal{F}$ such that $\mathbf{E}^F(r_n) = Kf$; *(IV)* r_n is a one-to-one transformation of some sufficient statistic. Moreover, $r_n \in \mathcal{F}$ is a Hilbert space-valued random variable (H-r.v.). We recall that, for a complete probability space $(Z, \mathcal{Z}, \mathbb{P})$, r_n is a H-r.v. if it defines a measurable map $r_n : (Z, \mathcal{Z}, \mathbb{P}) \to (\mathcal{F}, \mathfrak{B}_{\mathcal{F}})$, where $\mathfrak{B}_{\mathcal{F}}$ denotes the Borel σ-field generated by the open sets of $\mathcal{F}$.

Construction of r_n. Let $\mathfrak{T} \subseteq \mathbb{R}^l$, $l > 0$. To construct r_n we first select a function $k(t,x) : \mathfrak{T} \times S \to \mathbb{R}$ (or in $\mathbb{C}$) that is measurable in x for every $t \in \mathfrak{T}$ and that is non-constant in (t,x). The transformation r_n is then taken to be the expectation of $k(t,\cdot)$ under the empirical measure:

$$r_n(t) = \frac{1}{n}\sum_{i=1}^{n} k(t,x_i), \qquad \forall t \in \mathfrak{T}.$$

Define $\mathcal{F} = L^2(\mathfrak{T}, \mathfrak{B}_{\mathfrak{T}}, \rho)$ where ρ is a measure on $\mathfrak{T}$ and $\mathfrak{B}_{\mathfrak{T}}$ denotes the Borel σ-field generated by the open sets of $\mathfrak{T}$. The scalar product and norm on $\mathcal{F}$ are defined in the usual way and denoted by $\langle\cdot,\cdot\rangle$ and $\|\cdot\|$, respectively, with the same notation as for the inner product and norm in $\mathcal{E}$. The function $k(t,x)$ defines also a bounded operator $K : \mathcal{E} \to \mathcal{F}$ and must be such that, for every $\varphi \in \mathcal{E}$, $K\varphi \in \mathcal{F}$ and r_n is an H-r.v. with realizations in $\mathcal{F}$. Hence,

$$\begin{array}{rcl} K : \mathcal{E} & \to & \mathcal{F} \\ \varphi & \mapsto & \int k(t,x)\varphi(x)\Pi(dx). \end{array} \tag{2.3}$$

For every $f \in \mathcal{E}_M$, Kf is the expectation of $k(t,X)$ under F: $(Kf)(t) = \mathbf{E}^F(k(t,X))$. Under the true distribution F^* the expectation of r_n is Kf_* and the covariance function of r_n is: $\forall s,t \in \mathfrak{T}$,

$$\frac{1}{n}\sigma(t,s) = \mathbf{E}^* r_n(t) r_n(s) = \frac{1}{n}\left[\mathbf{E}^*\left(k(t,X)k(s,X)\right) - \mathbf{E}^*(k(t,X))\mathbf{E}^*(k(s,X))\right].$$

If the class of functions $\{k(t,\cdot), t \in \mathfrak{T}\}$ is Donsker then, as $n \to \infty$, the conditional distribution of $\sqrt{n}(r_n - Kf_*)$ weakly converges to a $\mathcal{GP}$ with covariance operator $\Sigma : \mathcal{F} \to \mathcal{F}$ defined as

$$\forall \psi \in \mathcal{F}, \qquad (\Sigma\psi)(t) = \int \sigma(t,s)\psi(s)\rho(ds) \tag{2.4}$$

which is one-to-one, linear, positive definite, self-adjoint and trace-class. In the following we assume that $\{k(t,\cdot), t \in \mathfrak{T}\}$ is Donsker such that r_n is approximately Gaussian: $r_n \sim \mathcal{GP}(Kf_*, \Sigma_n)$ where $\Sigma_n = \frac{1}{n}\Sigma$. Among all the functions $k(t,x)$ that satisfy the previous assumptions, one should keep only the ones such that the corresponding r_n is a one-to-one transformation of some sufficient statistic, as required in *(IV)* above. This, which will be tacitly assumed in the following of the paper, guarantees that the posterior distribution computed by using r_n does not depend on the particular choice of $k(t,x)$.

In our analysis we treat f_* as a realization of the random parameter f and Σ_n as known. Therefore, the sampling distribution of $r_n|f$ is $P^f = \mathcal{GP}(Kf, \Sigma_n)$ and we construct the posterior distribution based on it. In practice, Σ_n must be replaced by its empirical

counterpart. In finite sample, P^f is an approximation of the true sampling distribution but the approximation error vanishes as $n \to \infty$. Moreover, the approximating sampling distribution P^f is only used to construct the posterior distribution and the proofs of our asymptotic results do not rely on it. We give in the following two examples where r_n is a one-to-one transformation of sufficient statistics and K is injective.

Example 2.1 (Empirical cumulative distribution function (cdf))**.** Let $(X_1, \ldots, X_n)$ be an *i.i.d.* sample of $X \in \mathbb{R}$. A possible choice for $k(t,x)$ is $k(t,x) = 1\{x \leq t\}$, where $1\{A\}$ denotes the indicator function of the event A. In this case, $r_n(t) = F_n(t) := \frac{1}{n}\sum_{i=1}^n 1\{X_i \leq t\}$ is the empirical *cdf* and the operator K is $(K\varphi)(t) = \int_S 1\{s \leq t\}\varphi(s)\Pi(ds)$, $\forall \varphi \in \mathcal{E}$. By the Donsker's theorem, $F_n(\cdot)$ is asymptotically Gaussian with mean the true cdf $F_*(\cdot)$ and covariance operator characterized by the kernel: $\frac{1}{n}(F_*(s \wedge t) - F_*(s)F_*(t))$.

Example 2.2 (Empirical characteristic function)**.** Let $(x_1, \ldots, x_n)$ be an *i.i.d.* sample of $x \in \mathbb{R}$. Let $k(t,x) = e^{itx}$, so that $r_n(t) = c_n(t) := \frac{1}{n}\sum_{j=1}^n e^{itx_j}$ is the empirical characteristic function. In this case, the operator K is $(K\varphi)(t) = \int_S e^{its}\varphi(s)\Pi(ds)$, $\forall \varphi \in \mathcal{E}$. By the Donsker's theorem, $c_n(\cdot)$ is asymptotically a Gaussian process with mean the true characteristic function $c(\cdot) := \mathbf{E}^*[e^{itx}]$ and covariance operator characterized by the kernel: $\frac{1}{n}(c(s+t) - c(s)c(t))$.

The following lemma gives an useful characterization of the operator Σ in terms of K and its adjoint K^*. We recall that the adjoint K^* of a bounded and linear operator $K : \mathcal{E} \to \mathcal{F}$ is defined as the operator from $\mathcal{F}$ to $\mathcal{E}$ that satisfies $\langle K\varphi, \psi\rangle = \langle \varphi, K^*\psi\rangle$, $\forall \varphi \in \mathcal{E}$ and $\forall \psi \in \mathcal{F}$. In our case, an elementary computation shows that $(K^*\psi)(t) = \int_{\mathfrak{T}} k(t,x)\psi(t)\rho(dt)$, $\forall \psi \in \mathcal{F}$.

Lemma 2.2. *Let $K : \mathcal{E} \to \mathcal{F}$ be a bounded and linear operator defined as in* (2.3) *and $K^* : \mathcal{F} \to \mathcal{E}$ be its adjoint, that is, $(K^*\psi)(t) = \int_{\mathfrak{T}} k(t,x)\psi(t)\rho(dt)$, $\forall \psi \in \mathcal{F}$. The operator $\Sigma_n = \frac{1}{n}\Sigma$, with $\Sigma : \mathcal{F} \to \mathcal{F}$ defined in* (2.4) *takes the form*

$$\forall \psi \in \mathcal{F}, \quad \Sigma\psi = KM_fK^*\psi - (KM_f1)\langle M_f, K^*\psi\rangle \tag{2.5}$$

where $M_f : \mathcal{E} \to \mathcal{E}$ is the multiplication operator $M_f\varphi = f_*\varphi$, $\forall \varphi \in \mathcal{E}$.

We denote by $\mathfrak{D}$ the subset of $\mathcal{E}$ whose elements integrate to 0 with respect to Π:

$$\mathfrak{D} := \left\{ g \in \mathcal{E};\ \int_S g(x)\Pi(dx) = 0 \right\}$$

and denote by $K|_{\mathfrak{D}}$ the operator K restricted to $\mathfrak{D} \subset \mathcal{E}$. Remark that $\mathfrak{D}$ contains the functions in $\mathcal{E}$ that are the difference of *pdf*s of F with respect to Π. Moreover, $\mathcal{R}(\Omega_{0\theta}^{1/2}) \subset \mathfrak{D}$, where $\mathcal{R}(\cdot)$ denotes the range of an operator, and

$$\mathcal{R}(\Omega_{0\theta}^{1/2}) \subseteq \left\{ \varphi \in \mathcal{E};\ \int_S \varphi(h) h(\theta, x) \Pi(dx) = 0 \text{ and } \int_S \varphi(x) \Pi(dx) = 0 \right\}.$$

Remark that the equality holds when the sequence $(\lambda_j)_{j>d}$, used to construct $\Omega_{0\theta}$, is strictly positive. We also denote by $\mathfrak{D}_*$ the subset of $L^2(S, \mathfrak{B}_S, F_*)$ whose elements integrate to 0 with respect to F_*: $\mathfrak{D}_* := \left\{ g \in L^2(S, \mathfrak{B}_S, F_*);\ \int_S g(x) F_*(dx) = 0 \right\}$. It follows from Lemma 2.2 that we can write $\Sigma = HH^*$ where

$$\begin{array}{rccc} H: & \mathfrak{D}_* & \to & \mathcal{F} \\ & \varphi & \mapsto & Kf_*\varphi \end{array} \qquad \begin{array}{rccc} H^*: & \mathcal{F} & \to & \mathfrak{D}_* \\ & \psi & \mapsto & K^*\psi - \langle Kf_*, \psi \rangle. \end{array} \tag{2.6}$$

Remark that, H^* is the adjoint of H and is the projection of the adjoint of Kf_* (determined by using the scalar product in $L^2(S, \mathfrak{B}_S, F_*)$) onto $\mathfrak{D}_*$: $H^*\psi = \mathbb{P}K^*\psi$, where $\mathbb{P}\varphi = \varphi - \int_S \varphi(x) F_*(dx)$ is the projection operator of $L^2(S, \mathfrak{B}_S, F_*)$ on $\mathfrak{D}_*$.

2.3 Posterior distribution

The Bayesian model defines a joint distribution on (the Borel σ-field) of Λ and can be summarized in the following way:

$$\begin{aligned} \theta &\sim \mu(\theta) \\ f|\theta &\sim \mu(f|\theta) = \mathcal{GP}(f_{0\theta}, \Omega_{0\theta}), \quad \int h(\theta, x) f_{0\theta}(x) \Pi(dx) = 0 \quad \text{and} \quad \Omega_{0\theta}^{\frac{1}{2}}(1, h(\theta, \cdot)^T)^T = 0 \\ r_n|f, \theta &\sim r_n|f \sim P^f = \mathcal{GP}(Kf, \Sigma_n) \end{aligned} \tag{2.7}$$

where we use the $\mathcal{GP}$ approximation P^f. Theorem 1 in Florens and Simoni [2012] shows that the joint distribution of (f, r_n), conditional on θ, is:

$$\left(\begin{array}{c} f \\ r_n \end{array} \right) \Bigg| \theta \ \sim \ \mathcal{GP} \left(\left(\begin{array}{c} f_{0\theta} \\ Kf_{0\theta} \end{array} \right), \left(\begin{array}{cc} \Omega_{0\theta}, & \Omega_{0\theta}K^* \\ K\Omega_{0\theta}, & \Sigma_n + K\Omega_{0\theta}K^* \end{array} \right) \right) \tag{2.8}$$

where $(\Sigma_n + K\Omega_{0\theta}K^*) : \mathcal{F} \to \mathcal{F}$, $\Omega_{0\theta}K^* : \mathcal{F} \to \mathcal{E}$ and $K\Omega_{0\theta} : \mathcal{E} \to \mathcal{F}$. The marginal sampling distribution of r_n conditional on θ, obtained by integrating out f, is:

$$r_n|\theta \sim P_n^\theta := \mathcal{GP}(Kf_{0\theta}, \Sigma_n + K\Omega_{0\theta}K^*). \tag{2.9}$$

We now discuss the posterior distribution of the parameter of interest θ, denoted by $\mu(\theta|r_n)$. The conditional posterior distribution $\mu(f|r_n, \theta)$ of the nuisance parameter f, given θ, and its properties can be obtained by using results in Florens and Simoni [2014] and we briefly discuss it in Appendix **??**.

2.3.1 Posterior distribution of θ

The marginal posterior for θ, denoted by $\mu(\theta|r_n)$, is obtained by using the marginal sampling distribution P_n^θ given in (2.9). We first have to characterize the likelihood of P_n^θ with respect to an appropriate common dominating measure that will be denoted by P_n^0. The following theorem characterizes a probability measure P_n^0 which is equivalent to P_n^θ as well as the corresponding likelihood of P_n^θ with respect to P_n^0. Denote by $H^{-1} : \mathcal{R}(H) \to \mathfrak{D}_*$ (resp. $(H^*)^{-1} : \mathfrak{D}_* \cap \mathcal{R}(K^*) \to \mathcal{F}$) the left inverse of H (resp. right inverse of H^*) which takes the form $H^{-1}\cdot = f_*^{-1}K^{-1}\cdot$ (resp. $(H^*)^{-1}\cdot = (K^*)^{-1}\cdot$). Denote by $(l_{j\theta}, \rho_j(\theta), \psi_j(\theta))_{j\geq 0}$ the singular value decomposition of the operator $H^{-1}K\Omega_{0\theta}^{1/2}$ which is well defined by Lemma **??** in the Appendix if $K|_{\mathfrak{D}}$ is injective and f_* is bounded away from zero on S.

Theorem 2.1. *Let P_n^0 be a Gaussian measure with mean Kf_* and covariance operator $n^{-1}\Sigma$, i.e. $P_n^0 = \mathcal{GP}(Kf_*, n^{-1}\Sigma)$ with Σ defined in (2.4). For n fixed, if $K|_{\mathfrak{D}}$ is injective and f_* is bounded away from zero on S, then P_n^0 and P_n^θ are equivalent. Moreover, assume that $\forall j \geq 0$ and $\forall \theta \in \Theta$, $\psi_j(\theta) \in \mathfrak{D}_* \cap \mathcal{R}(K^*)$. Then, the Radon-Nikodym derivative is given by*

$$
\begin{aligned}
p_{n\theta}(r_n;\theta) &:= \frac{dP_n^\theta}{dP_n^0}(r_n) \qquad (2.10)\\
&= \prod_{j=0}^{\infty} \sqrt{\frac{n^{-1}}{n^{-1}+l_{j\theta}^2}} \exp\left\{ -\frac{1}{2}\sum_{j=0}^{\infty} \frac{\left(Z_j - \langle \sqrt{n}K(f_{0\theta}-f_*), (H^*)^{-1}\psi_j(\theta)\rangle\right)^2}{1+nl_{j\theta}^2} \right\} e^{\{\frac{1}{2}\|Z\|_\Sigma^2\}}
\end{aligned}
$$

where $Z := \sqrt{n}(r_n - Kf_*)$, $Z_j := \langle Z, (H^*)^{-1}\psi_j(\theta)\rangle$ for all $j \geq 0$, and $\|Z\|_\Sigma := \|\Sigma^{-1/2}Z\|$.

The quantity $\Sigma^{-1/2}$ denotes the inverse of the positive square root of Σ and $\|\Sigma^{-1/2}Z\|^2$ is defined as the limit in $\mathcal{F}$ of the series $\sum_{j=0}^m \sigma_j^{-2}\langle Z_j, \phi_j\rangle^2$ as $m \to \infty$ (where $\{\sigma_j^2, \phi_j\}_{j=0}^\infty$ is the eigensystem of Σ). By using (2.10), the (marginal) posterior distribution of θ takes the form (after simplifying the terms that do not depend on θ):

$$
\begin{aligned}
\mu(\theta|r_n) &= \frac{p_{n\theta}(r_n;\theta)\mu(\theta)}{\int_\Theta p_{n\theta}(r_n;\theta)\mu(\theta)d\theta} \qquad (2.11)\\
&= \frac{\prod_{j=0}^\infty \sqrt{\frac{1}{n^{-1}+l_{j\theta}^2}} \exp\left\{ -\frac{1}{2}\sum_{j=0}^\infty \frac{\left(Z_j - \langle \sqrt{n}K(f_{0\theta}-f_*), (H^*)^{-1}\psi_j(\theta)\rangle\right)^2}{1+nl_{j\theta}^2} \right\} \mu(\theta)}{\int_\Theta \prod_{j=0}^\infty \sqrt{\frac{1}{n^{-1}+l_{j\theta}^2}} \exp\left\{ -\frac{1}{2}\sum_{j=0}^\infty \frac{\left(Z_j - \langle \sqrt{n}K(f_{0\theta}-f_*), (H^*)^{-1}\psi_j(\theta)\rangle\right)^2}{1+nl_{j\theta}^2} \right\} \mu(\theta)d\theta}
\end{aligned}
$$

and can be used to compute a point estimator of θ. We propose to use the maximum a posterior (MAP) estimator θ_n defined as

$$
\begin{aligned}
\theta_n &:= \arg\max_{\theta\in\Theta} \mu(\theta|r_n) \\
&= \arg\max_{\theta\in\Theta} \prod_{j=0}^{\infty} \sqrt{\frac{1}{n^{-1}+l_{j\theta}^2}} \exp\left\{ -\frac{1}{2} \sum_{j=0}^{\infty} \frac{\left(Z_j - \langle \sqrt{n} K(f_{0\theta} - f_*), (H^*)^{-1}\psi_j(\theta)\rangle\right)^2}{1+nl_{j\theta}^2} \right\} \mu(\theta) \\
&= \arg\max_{\theta\in\Theta} \prod_{j=0}^{\infty} \sqrt{\frac{1}{n^{-1}+l_{j\theta}^2}} \exp\left\{ -\frac{1}{2} \sum_{j=0}^{\infty} \frac{\left(\langle \sqrt{n}(r_n - K f_{0\theta}), (H^*)^{-1}\psi_j(\theta)\rangle\right)^2}{1+nl_{j\theta}^2} \right\} \mu(\theta)
\end{aligned}
\tag{2.12}
$$

or the posterior mean estimator $\mathbf{E}(\theta|r_n) := \int_\Theta \theta\mu(\theta|r_n)d\theta$.

Remark 2.3. As discussed in Remark 2.2, when $d = p$ and θ can be written as a linear functional of f it is possible to use the different prior scheme described in Remark 2.2. With this different prior scheme, the posterior of θ is recovered from the $\mathcal{GP}$ posterior of f through the transformation $b(f)$.

2.3.2 Properties of the posterior distribution of θ

In this section we show two important results. The first one establishes that expression (2.10) is invariant to the choice of the measure Π used to define the nuisance parameter f and therefore the marginal posterior of θ is invariant to the choice of Π as well. More precisely the following proposition holds.

Proposition 2.1. *For a positive measure Π_1 on S, let $\mathcal{E}_{\Pi_1} := L^2(S, \mathfrak{B}_S, \Pi_1)$ and $\mathfrak{z} := \frac{d\Pi}{d\Pi_1}$. Let $\varphi : \mathcal{E} \to \mathcal{E}_{\Pi_1}$ be the transformation $\varphi(f) = f\mathfrak{z}$ and Φ be the set of measurable transformations defined as*

$$
\Phi := \left\{ \varphi : \mathcal{E} \to \mathcal{E}_{\Pi_1};\ \varphi(f) = f\mathfrak{z},\ \Pi_1 \text{ is a positive measure and } \sup_{x\in S} \frac{d\Pi_1(x)}{d\Pi(x)} < \infty \right\}.
$$

Then, the marginal posterior distribution $\mu(\theta|r_n)$ of θ is Φ-invariant.

This result shows that, once we integrate out the nuisance parameter f, the posterior distribution of θ is not affected by the choice of the dominating measure Π which only causes a transformation of the nuisance parameter.

Therefore, we can use two different dominating measures Π: one for the definition of the nuisance parameter f and one for the definition of $\mathcal{E}$, K and K^*. In particular, if $\sup_{x\in S} \frac{dF_*(x)}{d\Pi(x)} < \infty$ (where Π is used to define f and to construct its prior, and F_* is used to define $\mathcal{E}$, K and K^*) then, once we have specified the prior for the nuisance parameter $f = dF/d\Pi$, we deduce from it the prior of the transformation $\varphi(f) := fd\Pi/dF_*$ (see

the proof of Proposition 2.1). Therefore, the prior mean of $\varphi(f)$ is $f_{0\theta}d\Pi/dF_*$ and the prior covariance operator of $\varphi(f)$ writes in terms of an o.n.b. of $L^2(S,\mathfrak{B}_S,F_*)$ that we still denote by $\{\varphi_j\}_{j\geq 0}$ and where $\{\varphi_j\}_{j=1}^d$ are equal to the moment functions $h_j(\theta,x)$, $j=1,\ldots,d$, orthonormalized with respect to F_*: $\Omega_{0\theta}\cdot=\sum_{j>d}\lambda_j\langle\varphi_j,\cdot\rangle\varphi_j$. Moreover, the operator K and its adjoint are defined by using F_* instead of Π so that they are operators from (resp. to) the space $\mathcal{E}=L^2(S,\mathfrak{B}_S,F_*)$, and this does not change our inference on θ. Therefore, the operator H becomes $H\varphi=K|_{\mathfrak{D}_*}\varphi$, $\forall\varphi\in\mathfrak{D}_*$, and $H^{-1}\cdot=K^{-1}$. Because $H^{-1}K\Omega_{0\theta}K^*(H^*)^{-1}=\Omega_{0\theta}|_{\mathfrak{D}_*\cap\mathcal{R}(K^*)}$ it follows that the eigenfunctions $\{\psi_j(\theta)\}_{j\geq 0}$ are given by $\{\varphi_j\}_{j\geq 1}\in\mathfrak{D}_*\cap L^2(S,\mathfrak{B}_S,F_*)$ (which depend on θ and that are orthonormalized with respect to F_*) with corresponding eigenvalues $\{l_{j\theta}\}_{j\geq 0}=\{\lambda_j 1\{j>d\}\}_{j\geq 0}$ (which do not depend on θ). Remark that if K is injective, then $\mathcal{R}(K^*)=\overline{\mathcal{R}(K^*)}$ since $\mathfrak{N}(K)=\{0\}$ and K^* is a closed operator, and $\mathcal{R}(K^*)$ is dense in $\mathcal{E}$. Therefore, we can simplify the expression for $\mu(\theta|r_n)$ to:

$$\mu(\theta|r_n) \quad = \quad \frac{\exp\left\{-\frac{1}{2}\sum_{j=0}^{\infty}\frac{\langle\sqrt{n}(r_n-K\varphi(f_{0\theta})),(H^*)^{-1}\psi_j(\theta)\rangle^2}{1+n\lambda_j}\right\}\mu(\theta)}{\int_\Theta\exp\left\{-\frac{1}{2}\sum_{j=0}^{\infty}\frac{\langle\sqrt{n}(r_n-K\varphi(f_{0\theta})),(H^*)^{-1}\psi_j(\theta)\rangle^2}{1+n\lambda_j}\right\}\mu(\theta)d\theta}. \tag{2.13}$$

This simplification and the definition of K, K^* and $\mathcal{E}$ in terms of $\Pi=F_*$ must be understood in the following of the paper every time we explicitly assume $\sup_{x\in S}\frac{dF_*(x)}{d\Pi(x)}<\infty$. It is also important to remark that the orthonormalization of 1 and the moment functions $h_j(\theta,x)$, $j=1,\ldots,d$, with respect to F_* can be implemented either by using the Cholesky decomposition of (the empirical counterpart of) $\mathbf{E}^*[(1,h(\theta,X)^T)^T(1,h(\theta,X)^T)]$ or by using the Gram-Schmidt process.

The second result we are going to show[2] establishes a link between our Bayesian procedure, GEL estimators with quadratic criterion and the continuous updating GMM estimator. This relationship, given in Theorem 2.2 below, holds when the $\mathcal{GP}$ prior for $f|\theta$ is allowed to become diffuse. More precisely, let us rescale the prior covariance operator of $f|\theta$ by a positive scalar c so that the prior of $f|\theta$ may be written, for $f_{0\theta}\in\mathcal{E}_M$, as

$$\mu(f|\theta,c)\sim\mathcal{GP}(f_{0\theta},c\Omega_{0\theta}),\qquad\int h(\theta,x)f_{0\theta}(x)\Pi(dx)=0,\quad\Omega_{0\theta}^{1/2}(1,h(\theta,\cdot)^T)^T=0,\quad c\in\mathbb{R}_+.$$

Theorem 2.2. *Assume that* $\sup_{x\in S}\frac{dF_*(x)}{d\Pi(x)}<\infty$, $h_j(\theta,x)\in\mathcal{R}(K^*)$ *and* $\varphi_l\in\mathcal{R}(K^*)$, $\forall j=1,\ldots,d$, $l>d$ *and* $\forall\theta\in\Theta$, *and that* $\mathbf{E}^*[h(\theta,X_i)h(\theta,X_i)^T]$ *is nonsingular* $\forall\theta\in\Theta$. *Let* $\mu(f|\theta,c)\sim\mathcal{GP}(f_{0\theta},c\Omega_{0\theta})$, *with* $f_{0\theta}$ *and* $\Omega_{0\theta}$ *satisfying Restrictions 1 and 2, and* $c\in\mathbb{R}_+$. *Let* $\mu(\theta|r_n,c)$ *denote the (marginal) posterior of* θ *obtained by integrating out* f *from* P^f

[2]We thank Yuichi Kitamura for having suggested this research question.

with respect to $\mu(f|\theta,c)$*. Then,*

$$\lim_{c\to\infty}\mu(\theta|r_n,c)\propto\exp\left\{-\frac{1}{2}\left(\frac{1}{\sqrt{n}}\sum_{i=1}^{n}h(\theta,x_i)\right)^T V_n(\theta)^{-1}\left(\frac{1}{\sqrt{n}}\sum_{i=1}^{n}h(\theta,x_i)\right)\right\}\mu(\theta)$$

where $V_n(\theta):=\frac{1}{n}\sum_{i=1}^{n}h(\theta,x_i)h(\theta,x_i)^T$*.*

Remarks that in the theorem the limit $c\to\infty$ is taken after f has been marginalized out. The result in the theorem deserves some comments. First, it shows that, as the (conditional) prior on f becomes more and more diffuse, our marginal likelihood becomes the quasi-likelihood function (also called LIL) that has been used often in the literature, for instance by Chernozhukov and Hong [2003] and Kim [2002]. Therefore, the LIL naturally arises from a nonparametric Bayesian procedure, which places a $\mathcal{GP}$ prior on the set of functions in $\mathcal{E}$ that satisfy the EE, as the nonparametric prior becomes noninformative.

Second, Theorem 2.2 shows that, as the prior on f becomes noninformative, the MAP objective function is the same (up to constants and up to the prior) as the GEL objective function with quadratic criterion, see the proof of Theorem 2.1 in Newey and Smith [2004]. Moreover, as it can be deduced from Newey and Smith [2004, Theorem 2.1], the MAP objective function becomes a monotonic transformation of the continuous updating GMM objective function.

2.3.3 Properties of the MAP estimator of θ

By Proposition 2.1, our inference procedure is invariant to the choice of Π. Let us assume $\sup_{x\in S}\frac{dF_*(x)}{d\Pi(x)}<\infty$ and recall the discussion below Proposition 2.1 so that $\mathcal{E}=L^2(S,\mathfrak{B}_S,F_*)$. By Mercer's formula, under the assumption $\int_{\mathfrak{T}}\int_S|k(t,x)|^2F_*(dx)\rho(dt)<\infty$, we can write $k(t,x)=\sum_{j=1}^{\infty}\lambda_{jK}\phi_{j1}(t)\phi_{j2}(x)$ where $(\lambda_{jK},\phi_{j1},\phi_{j2})$ are such that $\int_S k(t,x)\phi_{j2}(x)F_*(dx)=\lambda_{jK}\phi_{j1}(t)$ and $\int_{\mathfrak{T}}k(t,x)\phi_{j1}(t)\rho(dt)=\lambda_{jK}\phi_{j2}(x)$, $\lambda_{jK}>0$ and $\{\phi_{j1}\}$, $\{\phi_{j2}\}$ are orthonormal sequences of $\mathcal{F}$ and $\mathcal{E}$, respectively. Therefore, $K\cdot=\sum_{j=1}^{\infty}\lambda_{jK}\phi_{j1}(t)\langle\phi_{j2},\cdot\rangle$ and since $\left((K^*)^{-1}\psi_j(\theta)\right)(t)=\sum_{j=1}^{\infty}\lambda_{jK}^{-1}\phi_{j1}(t)\langle\phi_{j2},\varphi_j\rangle$ under the assumptions of Theorem 2.1, we obtain: $\forall l\geq 0$

$$\begin{aligned}\langle\sqrt{n}r_n,(H^*)^{-1}\psi_l(\theta)\rangle&=\langle\sqrt{n}r_n,(K^*)^{-1}\varphi_l\rangle\\&=\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\sum_{j=1}^{\infty}\sum_{j'=1}^{\infty}\lambda_{jK}\langle\phi_{j1},\phi_{j'1}\rangle\phi_{j2}(x_i)\frac{1}{\lambda_{j'K}}\langle\phi_{j'2},\varphi_l\rangle\\&=\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\sum_{j=1}^{\infty}\phi_{j2}(x_i)\langle\phi_{j'2},\varphi_l\rangle=\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\varphi_l(x_i).\end{aligned}$$

It follows that the likelihood in (2.10) can be simplified and the MAP writes as:

$$\begin{aligned}
\theta_n &= \arg\max_{\theta\in\Theta} \mu(\theta|r_n) = \arg\max_{\theta\in\Theta} \left(\log p_{n\theta}(r_n;\theta) + \log \mu(\theta)\right) \\
&= \arg\max_{\theta\in\Theta} \Big(- \sum_{j=1}^{d} \langle \sqrt{n}(r_n - K\varphi(f_{0\theta})), (H^*)^{-1}\varphi_j\rangle^2 \\
&\qquad\qquad - \sum_{j>d} \langle \sqrt{n}(r_n - K\varphi(f_{0\theta})), (H^*)^{-1}\varphi_j\rangle^2 \frac{1}{1+n\lambda_j} + 2\log\mu(\theta)\Big)\frac{1}{2} \\
&= \arg\min_{\theta\in\Theta} \Big(\sum_{j=1}^{d} \left(\frac{1}{n}\sum_{i=1}^{n} \varphi_j(x_i) - \langle \varphi(f_{0\theta}), \varphi_j\rangle \right)^2 \\
&\qquad\qquad + \sum_{j>d} \left(\frac{1}{n}\sum_{i=1}^{n} \varphi_j(x_i) - \langle \varphi(f_{0\theta}), \varphi_j\rangle \right)^2 \frac{1}{1+n\lambda_j} - \log\mu(\theta)\Big) \qquad (2.14)
\end{aligned}$$

where we have eliminated the terms that do not depend on θ and we have used the fact that $\langle K\varphi(f_{0\theta}), (H^*)^{-1}\varphi_j\rangle = \langle \varphi(f_{0\theta}), \varphi_j\rangle$. Equation (2.14) is quite useful and allows to emphasize several aspects of our methodology.

I. The first term in (2.14) accounts for the EE. To see this, suppose that $(1, h(\theta,x)^T)^T$ have been orthonormalized by using the Gram-Schmidt process: $\varphi_0(x) = 1$ and $\varphi_j(x) = c_j(h_j(\theta,x) - \sum_{l=1}^{j-1} \mathbf{E}^*[h_l(\theta,X)h_j(\theta,X)]h_l(\theta,x) - \mathbf{E}^*[h_j(\theta,X)])$, $\forall 1 \le j \le d$ where c_j is the normalization constant. Then, $\langle \varphi(f_{0\theta}), \varphi_j\rangle = -c_j\mathbf{E}^*[h_j(\theta,X)]$, $\forall 1 \le j \le d$ and so $\frac{1}{n}\sum_{i=1}^{n}\varphi_j(x_i) - \langle \varphi(f_{0\theta}), \varphi_j\rangle = c_j(h_j(\theta,x) - \sum_{l=1}^{j-1}\mathbf{E}^*[h_l(\theta,X)h_j(\theta,X)]h_l(\theta,x))$ which corresponds to the orthonormalized moment functions with respect to F_*. Minimization of this term corresponds to the classical GMM.

II. The second term in (2.14) accounts for the extra information that we have, namely, the information contained in the subspace of $\mathcal{E}$ orthogonal to $span\{1, h_1(\theta,\cdot), \ldots, h_d(\theta,\cdot)\}$. This information, which is in general not exploited in EE frameworks (in frequentist as well as in Bayesian approaches), can be exploited thanks to the prior distribution and the prior mean $f_{0\theta}$ if the prior is not fixed but varies with n at an appropriate rate (see comment III below). On the contrary, if the prior is fixed then, as $n \to \infty$, the second term of (2.14) converges to 0 since $(1+n\lambda_j)^{-1} \to 0$. Remark also that $n^{-1}\sum_{i=1}^{n}\varphi_j(x_i) \to \mathbf{E}^*[\varphi_j(X)]$ a.s. and $\mathbf{E}^*[\varphi_j(X)] = 0$ because φ_j is orthogonal to 1 for $j > d$.

III. Expression (2.14) makes an explicit connection between the parametric case (infinite number of moment restrictions) and the semiparametric case (where only the first d moment conditions hold). The semiparametric case corresponds to the classical GEL or GMM approach while the parametric case corresponds to the maximum likelihood estimator (MLE). Indeed, the prior distribution for f specifies a parametric model for $f_{0\theta}$ which satisfies the d

moment restrictions and eventually other "extra" moment restrictions. The eigenvalues λ_j of the prior covariance operator play the role of weights of the "extra" moment restrictions and represent our "beliefs" concerning these restrictions. When we are very confident about these "extra" conditions, or equivalently we believe that $f_{0\theta}$ is close to f_*, then the λ_js are close to zero or converge to 0 faster than n^{-1} as $n \to \infty$. So, the prior distribution for f is degenerate on $f_{0\theta}$ (as n increases) when the parametric model is the true one. In that case, the MAP estimator will essentially be equivalent to the MLE that we would obtain if we use the prior mean function $f_{0\theta}$ as the likelihood. When we are very uncertain about $f_{0\theta}$ then the λ_js are very large and may tend to $+\infty$ (uninformative prior). In this case the MAP estimator will be close to the GMM estimator (up to a prior on θ).

2.3.4 Testing and moment selection procedures

Remark III in section 2.3.3 is important if one is interested in constructing testing procedures or doing moment selection. We are not going to develop a formal test/selection procedure here as this will make the object of a separated paper, but we would like to point out that our procedure suggests an easy way to test a parametric model against a semiparametric one characterized by a finite number of moment restrictions. We can deal with the two following situations where we assume $\sup_{x\in S} dF_*(x)/d\Pi(x) < \infty$:

[1.] We know that the distribution of the data F_* satisfies d moment restrictions and we want to test whether it has a particular parametric form. In this case, for a given probability distribution G such that $\int h(\theta, x)G(dx) = 0$ for a known vector of functions $h(\theta, x)$, the null hypothesis is $\mathcal{H}_0 : F_* = G$. An example is the univariate linear regression model: $Y = Z\theta + \varepsilon$, where f_* is the true joint *pdf* of $X := (Y, Z)^T$ and $\mathbf{E}^*(Y|Z) = Z\theta$ so that $h(\theta, X) = Y - Z\theta$. We may want to test that f_* belongs to a particular parametric class.

[2.] There are d EE of which we are sure and we want to test the validity of the other $D - d$ EE. The null hypothesis writes $\mathcal{H}_0 : \mathbf{E}^*(h_j(\theta, X)) = 0$ for $d < j \leq D$.

To treat [1.], we have to specify $f_{0\theta} = dG/d\Pi$. To treat [2.] we have to specify $f_{0\theta}$ such that it satisfies the extra EE: $\int_S h_j(\theta, x) f_{0\theta}(x)\Pi(dx) = 0$, $d < j \leq D$, and the $\{\varphi_j\}_{j=d+1}^D$ as the orthonormalized $h_j(\theta, x)$, $d < j \leq D$, with respect to F_* (so that the square term in the second term of (2.14) is zero only at the true θ_* if the extra EE are correctly specified). Then, for both the situations, the natural approach would be to treat the λ_js corresponding to the extra EE as hyperparameters for which a prior distribution is specified. The null hypothesis writes $\mathcal{H}_0 : \lambda_j = 0$ for all $j > d$ in [1.], and $\mathcal{H}_0 : \lambda_j = 0$ for all $d < j \leq D$ in [2.]. Then, the posterior distribution of λ_j may be used to draw a conclusion on the test: either by considering posterior odds ratio or by constructing encompassing tests.

To construct a prior for the λ_js let us write: $\lambda_j = c\rho_j$, $j > d$, where $c = tr\Omega_0$ and $\sum_{j=0}^{\infty} \rho_j = 1$. The $\{\lambda_j\}_{j=1}^{d}$ corresponding to the first d EE do not affect the trace of Ω_0 since they are equal to 0. We propose two alternatives priors.

Dirichlet prior. Suppose that we want to test $\mathcal{H}_0 : \lambda_j = 0$ for all $d < j \leq D$. Then one may specify a Dirichlet prior for $(\rho_{d+1}, \ldots, \rho_{D-1})$:

$$\mu_\rho(\rho_{d+1}, \ldots, \rho_{D-1}|\nu) \propto \prod_{j=d+1}^{D-1} \rho_j^{\nu_j - 1} \left(1 - \sum_{j=d+1}^{D-1} \rho_j\right)^{\nu_D - 1} \prod_{j=d+1}^{D-1} I(\rho_j \geq 0) I\left(\sum_{j=1}^{D-1} \rho_j \leq 1\right)$$

where $\nu = (\nu_{d+1}, \ldots, \nu_D)$.

Prior on $c > 0$. Suppose that we want to test $\mathcal{H}_0 : \lambda_j = 0$ for all $j > d$. Thus, the null hypothesis may be written $\mathcal{H}_0 : c = 0$ and a prior for c may be any distribution with support contained in the positive real semi-axis, for example an inverse gamma distribution.

3 Asymptotic Analysis

In this section we focus on the frequentist asymptotic properties of our approach for $n \to \infty$. For this analysis we use the true probability measure P^* which corresponds to the true DGP F_* and maintain the assumption that $\sup_{x \in S} \frac{dF_*(x)}{d\Pi(x)} < \infty$ where Π is used to define the nuisance parameter f and F_* to define K, K^* and $\mathcal{E}$ (see the discussion below Proposition 2.1). Therefore, $\varphi(f) = f_*^{-1} f$. We show: *(i)* consistency of the posterior of θ (Theorem 3.1), *(ii)* frequentist consistency of the MAP estimator θ_n (Theorem 3.2), and *(iii)* convergence in Total Variation of $\mu(\theta|r_n)$ towards a normal distribution (section 3.2). In the following, for every $\tilde{\theta} \in \Theta$ and $\delta > 0$ we denote by $B(\tilde{\theta}, \delta)$ the closed ball centered in $\tilde{\theta}$ with radius δ, that is, $B(\tilde{\theta}, \delta) = \{\theta \in \Theta;\ \|\theta - \tilde{\theta}\| \leq \delta\}$, where here $\|\cdot\|$ denotes the Euclidean norm in $\mathbb{R}^p$. Moreover, denote $\delta_n = n^{-1/2}$.

3.1 Posterior Consistency

We first state the following assumptions.

A1. The true parameter θ_* belongs to the interior of a compact convex subset Θ of $\mathbb{R}^d$ and is the unique solution of $\mathbf{E}^*[h(\theta, X)] = 0$.

A2. The singular functions $\{\psi_j(\theta), \rho_j(\theta)\}$ are continuous functions of θ.

A3. The prior mean function $f_{0\theta}$ and the prior distribution $\mu(\theta)$ are continuous in θ.

A4. For every j and every n, $\mathbf{E}^* \langle \sqrt{n}(r_n - K\varphi(f_*)), (H^*)^{-1}\psi_j(\theta) \rangle^4 < \infty$.

Assumption A1 is a standard assumption in the literature on EE. Assumptions A2 and A3 can be easily satisfied since $f_{0\theta}$ and the operators $\Omega_{0\theta}$, K and Σ are chosen by the statistician. The next theorem gives concentration of the posterior distribution around the true value θ_*.

Theorem 3.1. *Let the assumptions of Theorem 2.1 and Assumptions A1-A4 be satisfied and assume* $\sup_{x\in S} \frac{dF_*(x)}{d\Pi(x)} < \infty$. *Then, for any prior* $\mu(\theta)$ *that puts enough mass in a neighborhood of* θ_* *and any sequence* $M_n \to \infty$,

$$\mu\left(\sqrt{n}\|\theta - \theta_*\| > M_n | r_n\right) \to 0 \tag{3.1}$$

in P^**-probability as* $n \to \infty$.

Given the result of the theorem, the next result follows.

Theorem 3.2. *Let the assumptions of Theorem 3.1 be satisfied. Then,* $\theta_n \xrightarrow{p} \theta_*$ *in* P^**-probability as* $n \to \infty$.

3.2 Asymptotic Normality

In this section we establish asymptotic normality of $\mu(\theta|r_n)$ for the Bayesian model described in (2.7). This result applies to the case $d > p$ (which is our main interest) as well as to the case $d = p$. In appendix **??** we establish, under different assumptions, asymptotic normality of $\mu(\theta|r_n)$ for the case $d = p$ described in Remarks 2.2 and 2.3 where the prior for θ is deduced from the prior for f.

For some $\tau \in \mathbb{R}^p$, let

$$s_n(\tau) := p_{n,\theta_*+\delta_n\tau}(r_n; \theta_* + \delta_n\tau).$$

We assume that there exist a random vector $\tilde{\ell}_*$ and a nonsingular matrix $\tilde{I}_*$ (that depend on the true θ_* and f_*) such that the sequence $\tilde{\ell}_*$ is bounded in probability, and satisfy

$$\log \frac{s_n(\tau)}{s_n(0)} = \frac{1}{\sqrt{n}}\tau^T \widetilde{I}_* \widetilde{\ell}_* - \frac{1}{2}\tau^T \widetilde{I}_* \tau + o_p(1) \tag{3.2}$$

for every random sequence τ which is bounded in P^*-probability. Condition (3.2) is known as the integral local asymptotic normality assumption which is used to prove asymptotic normality of semiparametric Bayes procedures, see e.g. Bickel and Kleijn [2012]. In Appendix **??** we prove that, if $\sup_{x\in S} \frac{dF_*(x)}{d\Pi(x)} < \infty$, then equation (3.2) holds with

$$\widetilde{I}_* = -\mathbf{E}^* \left[\frac{\partial h(\theta_*, x)}{\partial \theta}\right] \left[\mathbf{E}^* h(\theta_*, X) h(\theta_*, X)^T\right]^{-1} \mathbf{E}^* \left[\frac{\partial h(\theta_*, X)}{\partial \theta^T}\right]$$

if $\mathbf{E}^*\left[h(\theta_*, X)h(\theta_*, X)^T\right]$ is nonsingular. For two probability measures P_1 and P_2 absolutely continuous with respect to a positive measure Q, define the total variation (TV) distance as $||P_1 - P_2||_{TV} := \int |f_1 - f_2| dQ$ where f_1 and f_2 are the Radon-Nikodym derivatives of P_1 and P_2, respectively, with respect to Q. The following theorem shows that under (3.2) the posterior distribution of $\sqrt{n}(\theta - \theta_*)$ converges in the TV distance to a Normal distribution with mean $\Delta_* := \frac{1}{\sqrt{n}}\widetilde{\ell}_*$ and variance $\widetilde{I}_*^{-1}$.

Theorem 3.3. *Let the assumptions of Theorem 2.1 be satisfied. Assume that A1-A3,* (3.1) *and* (3.2) *hold and that the prior $\mu(\theta)$ puts enough mass in a neighborhood of θ_*. If $\mu(\sqrt{n}(\theta - \theta_*)|r_n)$ denotes the posterior of $\sqrt{n}(\theta - \theta_*)$, then:*

$$\|\mu(\sqrt{n}(\theta - \theta_*)|r_n) - \mathcal{N}(\Delta_*, \widetilde{I}_*^{-1})\|_{TV} \to 0 \tag{3.3}$$

in P^*-probability as $n \to \infty$.

4 Implementation

In this section we first explain the numerical implementation of our procedure and then show the results of three simulations. Let M be the number of discretization points and $x_{eval} := (x_1, \ldots, x_M)$ (resp. $t_{eval} := (t_1, \ldots, t_M)$) be (equidistant) discretization points for functions in $\mathcal{E}$ (resp. in $\mathcal{F}$). The procedure consists of Markov Chain Monte Carlo (MCMC) draws from the posterior (2.11) based on the Metropolis-Hastings (M-H) algorithm, see for instance Robert [2002]. Beforehand, the researcher must chose: *(a)* the measures $\Pi(x)$ and $\rho(t)$, *(b)* the function $k(t, x)$ (for instance: $k(t, x) = 1\{x \leq t\}$, $k(x, t) = e^{itx}$, or $k(x, t) = e^{tx}$), *(c)* the sequence λ_j (for instance: $\lambda_j = j^{-\alpha}$ with $\alpha > 1$, or α^{-j} with $\alpha > 1$), *(d)* the prior $\mu(\theta)$ and *(e)* the auxiliary distribution, say $g(\theta)$ for the M-H. To construct $\Omega_{0\theta}$ the series is truncated at a value J which is chosen such that λ_J is small (for instance of the order 10^{-3}). In the following, functions and operators are replaced by their discretized versions, denoted with a subindex $\mathfrak{d}$.

Before starting the M-H algorithm, compute the following quantities:

I. Compute r_n and discretize it to get the M-vector $r_{n,\mathfrak{d}} := (r_n(t_i))_{i=1,\ldots,M}$.

II. Compute $K_\mathfrak{d}$ as the $M \times M$ matrix with generic element $(k(t_i, x_j)\Pi(x_j))_{i,j=1,\ldots,M}$.

III. Compute $\Sigma_\mathfrak{d}$ as the $M \times M$ matrix with generic element $(\sigma(t_i, t_l)\rho(t_l))_{i,l=1,\ldots,M}$ where $\sigma(t_i, t_l) = n^{-1}\sum_{j=1}^n (k(t_i, x_j)k(t_l, x_j)) - r_n(t_i)r_n(t_l)$. Due to the discretization, the matrix $\Sigma_\mathfrak{d}$ could be ill-conditioned in some cases and then has to be regularized.

Then, implement the M-H algorithm where for every value $\tilde{\theta}^{(j)}$ drawn from the auxiliary distribution $g(\theta)$, the expression of $\log p_{n\theta}(r_n; \theta)$ in (2.10) is evaluated through the following

steps:

1. Evaluate the vector function $h(\theta, x)$ at $\tilde{\theta}^{(j)}$ and the M discretization points x_{eval} so to have a $(d \times M)$ matrix $h_{\mathfrak{d}} := (h(\tilde{\theta}^{(j)}, x_1), \ldots, h(\tilde{\theta}^{(j)}, x_M))$. Let ι_M denotes the M dimensional row-vector of ones.

2. Orthonormalize the set of vectors given by ι_M and the rows of $h_{\mathfrak{d}}$ (by using *e.g.* the Gram-Schmidt process) to obtain the $d+1$ M-vectors $\varphi_{j,\mathfrak{d}} := (\varphi_j(x_1), \ldots, \varphi_j(x_M))^T$ for $0 \leq j \leq d$.

3. Complete the basis by first constructing $J-d-1$ vectors ι_M and then orthonormalizing them sequentially by using *e.g.* the Gram-Schmidt process to obtain the remaining components of the basis $\varphi_{j,\mathfrak{d}}$, $j = d+1, \ldots, J$. Denote $\varphi_M := (\varphi_j(x_i))_{\substack{i=1,\ldots,M \\ j=0,\ldots,J-1}}$ and $\varphi_\Pi := (\varphi_j(x_i)\Pi(x_i)step)_{\substack{i=1,\ldots,M \\ j=0,\ldots,J-1}}$, which have both dimension $J \times M$ and where *step* denotes the discretization step.

4. Construct $\Omega_{0\theta,\mathfrak{d}}$ as $\Omega_{0\theta,\mathfrak{d}} = \varphi_M^T diag(\lambda_j)_{j=0,\ldots,J}\varphi_\Pi$, where $\lambda_j = 0$ for $j = 0, \ldots, d$.

5. Select a prior mean function $f_{0\theta}$ and discretize it to get an M-vector $f_{0\theta,\mathfrak{d}} := (f_{0\theta}(x_i))_{i=1,\ldots,M}$.

6. Compute the singular value decomposition $(l_{j\theta}, \rho_j(x_i), \psi_j(t_i))_{\substack{j=0,\ldots,J-1 \\ i=1,\ldots,M}}$ of the matrix $(\sqrt{\Sigma_{\mathfrak{d}}})^{-1} K_{\mathfrak{d}} \sqrt{\Omega_{0\theta,\mathfrak{d}}}$, where for a positive definite matrix A, $\sqrt{A}$ denotes its positive square root. Denote $\psi_{j,\mathfrak{d}} := (\psi_j(t_1), \ldots, \psi_j(t_M))^T$.

7. Compute the terms in $\log p_{n\theta}(r_n; \theta)$ that depend on θ by discretizing the scalar product:

$$\frac{1}{2}\sum_{j=1}^{J} \log\left(1 + n l_j^2\right) + \frac{1}{2}\sum_{j=1}^{J} \frac{\left((r_{n,\mathfrak{d}} - K_{\mathfrak{d}} f_{0\theta,\mathfrak{d}})^T diag(\rho(t_i))_{i=1,\ldots,M} (\sqrt{\Sigma_{\mathfrak{d}}})^{-1} \psi_{j,\mathfrak{d}}\right)^2}{1 + n l_{j\theta}^2}.$$

We remark that the Gram-Schmidt orthonormalization in steps 2 and 3 is extremely simple and fast to implement even for complicated moment functions. By using the decomposition $\Sigma = HH^*$ with H, H^* as described in (2.6), the previous steps can be simplified in the following way: steps III and 6 do not have to be implemented and $(\sqrt{\Sigma_{\mathfrak{d}}})^{-1}$ in step 7 has to be replaced by $(K^*)^{-1}$. Moreover, $\psi_{j,\mathfrak{d}}$ and l_j^2 have to be replaced by $(\varphi_j(x_i))_{i=1,\ldots,M}$ and λ_j, respectively. We now show the finite sample properties of our $\mathcal{GP}$ approach with the help of three simulations. The interest in using a $\mathcal{GP}$ prior will be made evident in the more complicated examples where there are overidentifying restrictions.

4.1 Just identification and prior on θ through $\mu(f)$

Let the parameter θ of interest be the population mean: $\theta = \int xf(x)dx$ and $h(\theta, x) = (\theta - x)$. This example considers the $d = p$ case where θ is a linear functional of f and the prior of θ is deduced from the prior of f, denoted by $\mu(f)$. Suppose that the support of F_* is $\mathbb{R}$, so that θ can take every value in $\mathbb{R}$. Then the prior $\mu(f)$ is a $\mathcal{GP}$ which is unrestricted except for the fact that it must generate trajectories that integrate to 1 a.s., namely, $\mu(f) \sim \mathcal{GP}(f_0, \Omega_0)$ where f_0 is a *pdf* and Ω_0 is such that $\Omega_0^{1/2} 1 = 0$. Therefore, the prior distribution of θ is Gaussian with mean $\langle f_0, \iota \rangle$ and variance $\langle \Omega_0 \iota, \iota \rangle$ where ι denotes the identity functional, that is, $\iota(x) = x$. The posterior distribution of θ is as in (**??**) with $g = \iota$ and $A = \Omega_0 K^* \left(n^{-1}\Sigma + K\Omega_0 K^*\right)^{-1}$. We illustrate how to construct in practice the covariance operator Ω_0. Let $S = \mathbb{R}$; the Hermite polynomials $\{H_j\}_{j\geq 0}$ form an orthogonal basis of $L^2(\mathbb{R}, \mathfrak{B}, \Pi)$ for $d\Pi(x) = e^{-x^2/2}dx$ and can be used to construct the eigenfunctions of Ω_0. The first few Hermite polynomials are $\{1, x, x^2 - 1, (x^3 - 3x), \ldots\}$ and they are orthogonal with respect to Π: $\int_{\mathbb{R}} H_l(x)H_j(x)e^{-x^2/2}dx = \sqrt{2\pi}n!\delta_{lj}$, where δ_{lj} is equal to 1 if $l = j$ and to 0 otherwise. The operator Ω_0 is constructed as

$$\Omega_0 \cdot \quad = \quad \sigma_0 \sum_{j=0}^{\infty} \lambda_j \frac{1}{\sqrt{2\pi}n!} \langle H_j, \cdot \rangle H_j$$

where $H_{j+1}(x) = xH_j(x) - jH_{j-1}(x)$, $\lambda_0 = 0$ and $\{\lambda_j,\ j \geq 1\} = \{a^j,\ j \geq 1\}$ with $a < 1$.

In our simulation exercise we generate $n = 1000$ *i.i.d.* observations $(x_1, \ldots, x_n)$ from a $\mathcal{N}(1, 1)$ distribution and construct the function $r_n = n^{-1} \sum_{i=1}^n e^{tx_i}$ as the empirical Laplace transform. Therefore, $f_*(x) = \frac{1}{\sqrt{2\pi}} e^{-(1-2x)/2}$ and $\theta_* = 1$. We set $\mathfrak{T} = \mathbb{R}$ and $\rho = \Pi$. Thus, the operators K and K^* take the form

$$\forall \phi \in \mathcal{E},\ \ K\phi = \int_{\mathbb{R}} e^{tx}\phi(x)e^{-x^2/2}dx \quad \text{ and } \forall \psi \in \mathcal{F}, \qquad K^*\psi = \int_{\mathbb{R}} e^{tx}\psi(t)e^{-t^2/2}dt.$$

The prior mean function f_0 is set equal to the Π-pdf of a $\mathcal{N}(\varrho, 1)$ distribution: $f_0(x) = \frac{1}{\sqrt{2\pi}} e^{-(\varrho^2 - 2\varrho x)/2}$, $\varrho = 2$. We show in Figure 1 the prior and posterior distribution of θ. We also show the prior mean (magenta asterisk), the posterior mean (blue asterisk) and the MAP (red asterisk) of θ. The posterior mean of θ is computed by discretizing the inner product $\langle \mathbf{E}(f|r_n), \iota \rangle$. The pictures are obtained for $a = 0.3$ and $\sigma_0 = 1$. The number of discretization points, used to approximate the integrals, is equal to 1000.

4.2 Just identification and prior on θ

We consider the same framework as in the previous example where the parameter θ of interest is the population mean: $\theta = \int xf(x)dx$ and $h(\theta, x) = (\theta - x)$, but now we specify

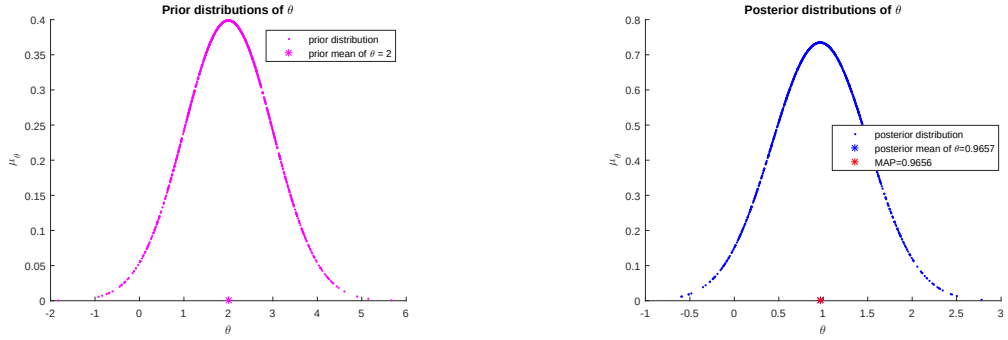

Figure 1: Prior and posterior distributions and means of θ. The true value of θ is $\theta_* = 1$.

a joint proper prior distribution on (θ, f). We specify a marginal prior $\mu(\theta)$ on θ and a conditional prior on f given θ. While $\mu(\theta)$ can be arbitrarily chosen, $\mu(f|\theta)$ is specified as a $\mathcal{GP}$ constrained to generate functions that integrate to 1 and that have mean equal to θ a.s., as described in section 2.1.

Compared to the approach in section 4.1, this approach allows to easily incorporate any prior information that one may have about θ. In fact, incorporating the information on θ through the prior distribution on f is complicated while to incorporate such an information directly in the prior distribution of θ results to be very simple. In particular, the approach of this section works even when θ takes values in a compact subset of $\mathbb{R}^p$, while the approach of section 4.1 does not work in this case.

Let us suppose that $m = 1$, $S = [-1, 1]$ and let Π and ρ be the Lebesgue measure. Then, the covariance operator $\Omega_{0\theta}$ can be constructed by using Legendre polynomials since the second Legendre polynomial $P_1(x) = x$ allows to implement the constraint on θ. Because the moment function is separable in θ and x, the prior covariance operator does not depend on θ, so that we denote it by Ω_0. The first few Legendre polynomials are $\{1, x, (3x^2 - 1)/2, (5x^3 - 3x)/2, \ldots\}$ and they are orthogonal with respect to the L^2 inner product on $[-1, 1]$: $\int_{-1}^{1} P_l(x)P_j(x)dx = 2/(2j+1)\delta_{lj}$, where δ_{lj} is equal to 1 if $l = j$ and to 0 otherwise. Moreover, the Legendre polynomial obey the recurrence relation $(j + 1)P_{j+1}(x) = (2j + 1)xP_j(x) - jP_{j-1}(x)$ which is useful for computing Ω_0 in practice. The normalized Legendre polynomials form a basis for $L^2[-1, 1]$ so that we can construct the operator Ω_0 as

$$\Omega_0 \cdot = \sigma_0 \sum_{j=2}^{\infty} \lambda_j \frac{2j+1}{2} \langle P_j, \cdot \rangle P_j$$

where we have set $\lambda_0 = \lambda_1 = 0$ in order to implement the constraints. The remaining λ_j, $j \geq 2$ can be chosen in an arbitrary way provided that $\sum_{j \geq 2} \lambda_j < \infty$. The constant $\sigma_0 > 0$ can be set to an arbitrary value and has the purpose of tuning the size of the prior covariance.

Many orthogonal polynomials are suitable for the construction of $\Omega_{0\theta}$ and they may be used to treat cases where S is different from $[-1,1]$.

We perform two simulations exercises: the first one makes use of the empirical cdf to construct r_n: $r_n(t) = F_n(t) := n^{-1}\sum_{i=1}^n 1\{x_i \leq t\}$ and the second one uses the empirical moment generating function $r_n(t) = n^{-1}\sum_{i=1}^n e^{tx_i}$. In both the simulations we use Legendre polynomials and we generate $n = 1000$ *i.i.d.* observations $(x_1, \ldots, x_n)$ from a $\mathcal{N}(0,1)$ distribution truncated to the interval $[-1,1]$. The prior distribution for θ is uniform over the interval $[-1,1]$. The prior mean function $f_{0\theta}$ is taken equal to the *pdf* of a Beta distribution with parameters p_θ and q and with support $[-1,1]$:

$$f_{0\theta}(x) = \frac{(x+1)^{p_\theta - 1}(1-x)^{q-1}}{B(p_\theta, q)2^{p_\theta+q-1}}. \tag{4.1}$$

We use the notation p_θ to stress the dependence on θ of this shape parameter. We fix $q = 2$ and recover p_θ such that $\int_{-1}^1 x f_{0\theta}(x)dx = \theta$. It is easy to see that for our Beta distribution: $\int_{-1}^1 x f_{0\theta}(x)dx = \frac{p_\theta - q}{p_\theta + q}$. The covariance operator $\Omega_{0\theta}$ is constructed by using the Legendre polynomials, $\lambda_j = j^{-1.7}$ and $\sigma_0 = 5$.

Since the posterior distribution $\mu(\theta|r_n)$ can not be computed in a closed-form we simulate from it by using a M-H algorithm. To implement this algorithm we have to select an auxiliary *pdf* $g(\theta)$. We summarize the simulation scheme.

1. Draw a n *i.i.d.* sample $(x_1, \ldots, x_n)$ from f_* (where f_* is a $\mathcal{N}(0,1)$ truncated to $[-1,1]$);
2. compute $r_n = F_n$ or $r_n = n^{-1}\sum_{i=1}^n e^{tx_i}$;
3. draw $\theta \sim \mathcal{U}[-1,1]$ and denote it $\tilde{\theta}$;
4. compute p_θ as $p_\theta = \frac{(\tilde{\theta}+1)2}{1-\tilde{\theta}}$ (where we have fixed $q = 2$);
5. compute $f_{0\theta}$ as in (4.1) with parameters $(p_\theta, q = 2)$ and discretize it;
6. draw θ from the marginal posterior distribution of θ by using a M-H algorithm with the following auxiliary *pdf* (triangular distribution):

$$g(\xi;\theta) = \frac{\xi+1}{\theta+1}1\{\xi \in [-1,\theta)\} + \frac{1-\xi}{1-\theta}1\{\xi \in [\theta,1]\}$$

as described above. We draw 10000 values and discard the first 5000. The initial value for the algorithm is $\theta = 0.5$.

We represent in Figure 2a the results for the simulation with $r_n(t) = F_n(t)$ and in Figure 2b the results for $r_n = n^{-1}\sum_{i=1}^n e^{tx_i}$: the blue asterisk represents the posterior

mean estimate while the red asterisk represents the MAP estimate. These figures also show the marginal posterior distribution of θ (dashed blue line) approximated by using a kernel smoothing and 5000 drawings from the posterior. In both the simulations, $n = 1000$ and the number of discretization points, used to approximate the integrals, is equal to 1000.

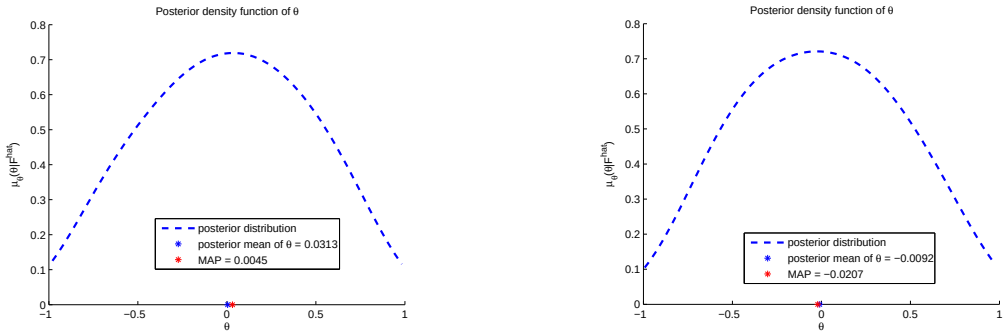

(a) $r_n(t) = F_n(t)$ (b) $r_n = n^{-1} \sum_{i=1}^{n} e^{tx_i}$

Figure 2: Estimations of θ based on the posterior distribution: posterior mean and MAP. The true value of θ is $\theta_* = 0$ and $n = 1000$.

4.3 Overidentified case

Let us consider the case in which x is univariate and the one-dimensional parameter of interest θ is characterized by the moment conditions $\mathbf{E}^F(h(\theta, X)) = 0$ with $h(\theta, x) = (x - \theta, 2\theta^2 - x^2)^T$. For instance, this arises when the true DGP F_* is an exponential distribution with parameter θ. The prior $\mu(\theta)$ is specified as a $\mathcal{U}[\theta_* - 1, \theta_* + 1]$.

The moment conditions are incorporated in the prior $\mu(f|\theta)$ for f as described in section 2.1. We chose $\Pi(dx) = e^{-x}dx$ and $r_n(t) = F_n(t)$. We first orthonormalize the moment functions 1, $x-\theta$, $2\theta^2-x^2$ with respect to Π and then complete the bases by using the Gram-Schmidt orthonormalization process. The inner products in $\mathcal{E}$ are approximating by using the trapezoidal rule on equally spaced subintervals of the interval $[\min x_i - 1, \max x_i + 1]$. We use polynomially decreasing eigenvalues for $\Omega_{0\theta}$: $\lambda_j = j^{-1.7}$. Finally, to construct $\Omega_{0\theta}$ we truncate the series at $J = 300$ since after that the value of λ_j is of the order 10^{-5} and then can be considered zero.

In our simulation, we generate $n = 500$ observations $x_1, \ldots, x_n$ from an exponential distribution with parameter $\theta_* = 2$. The measure $\rho(dt)$, necessary to construct K^*, is taken equal to the Lebesgue measure. Operators K and K^* are approximated by using the trapezoidal rule on equally spaced subintervals of the following intervals: $[\min x_i - 1, \max x_i + 1]$ for K and $[\min x_i, \max x_i]$ for K^*. The operator Σ is approximated in a similar way. Because of this discretization, the operator Σ is ill-conditioned and hence we regularize it by adding to it the identity matrix scaled by n^{-1}.

The prior mean function $f_{0\theta}$ is chosen by using a two-step procedure where in the first step we compute $\tilde{f} = (0.1I + K^*K)^{-1}K^*r_n$ and in the second step we project it on $\Lambda(\theta)$ for a given θ. Alternatively, $f_{0\theta}$ can be constructed a described before Remark 2.2.

To draw from the posterior distribution of θ, we use a M-H algorithm. To implement this algorithm we use, as auxiliary distribution, a $\chi^2_{\lceil\theta\rceil}$ distribution. The posterior distribution, its mean and its mode obtained in this simulation are plotted in Figure 3a. The posterior density function has been obtained by kernel smoothing with a Gaussian kernel and a bandwidth equal to 0.3.

Finally, we have repeated the same Monte Carlo simulation 100 times and have computed the average of the posterior mean estimators and MAP estimators. We report the results in Figure 3b together with the posterior density, mean and MAP obtained in each simulation.

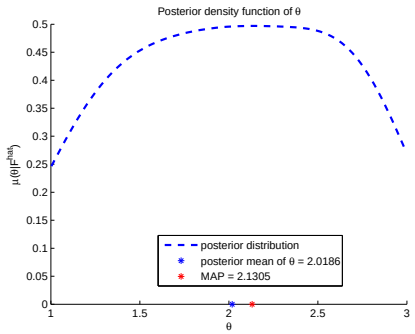

(a) Result for one simulation.

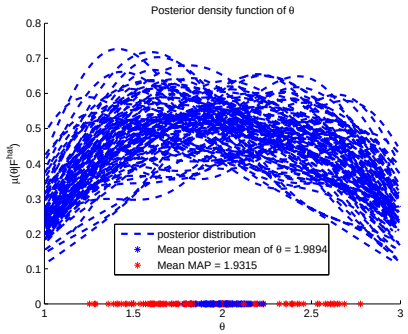

(b) Average of posterior means and MAP estimators over 100 Monte Carlo iterations.

Figure 3: Overidentified case $d > p$. Posterior distributions of θ, mean and MAP estimators. $r_n = F_n$ and the true value of θ is $\theta_* = 2$.

References

P. J. Bickel and B. J. K. Kleijn. The semiparametric bernstein–von mises theorem. *The Annals of Statistics*, 40(1):206–237, 2012.

L. Bornn, N. Shephard, and R. Solgi. Moment conditions and bayesian nonparametrics. Technical report, arXiv:1507.08645, 2015.

M. Broniatowski and A. Keziou. Divergences and duality for estimation and test under moment condition models. *Journal of Statistical Planning and Inference*, 142(9):2554 – 2573, 2012.

G. Chamberlain and G. W. Imbens. Nonparametric applications of bayesian inference. *Journal of Business & Economic Statistics*, 21(1):12–18, 2003.

S. Chaudhuri, D. Mondal, and T. Yin. Hamiltonian monte carlo sampling in bayesian empirical likelihood computation. *Journal of the Royal Statistical Society, Series B*, forthcoming, 2017.

S. X. Chen and I. Van Keilegom. A review on Empirical Likelihood methods for regression. *TEST*, 18(3):415–447, 2009.

V. Chernozhukov and H. Hong. An mcmc approach to classical estimation. *Journal of Econometrics*, 115(2):293–346, 2003.

S. Chib, M. Shin, and A. Simoni. Bayesian empirical likelihood estimation and comparison of moment condition models. Technical report, arXiv:1606.02931, 2016.

J.-P. Florens and A. Simoni. Regularized posteriors in linear ill-posed inverse problems. *Scandinavian Journal of Statistics*, 39(2):214–235, 2012.

J.-P. Florens and A. Simoni. Regularizing priors for linear inverse problems. *Econometric Theory*, FirstView, 11 2014.

M. Grendar and G. Judge. Asymptotic equivalence of empirical likelihood and bayesian map. *The Annals of Statistics*, 37(5A):2445–2457, 2009.

L. P. Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 50:1029–1054, 1982.

L. P. Hansen and K. J. Singleton. Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica*, 50(5):1269–1286, 1982.

L. P. Hansen, J. Heaton, and A. Yaron. Finite-sample properties of some alternative gmm estimators. *Journal of Business & Economic Statistics*, 14(3):262–280, 1996.

G. W. Imbens. One-step estimators for over-identified generalized method of moments models. *The Review of Economic Studies*.

J.-Y. Kim. Limited information likelihood and bayesian analysis. *Journal of Econometrics*, 107(1-2):175–193, 2002.

Y. Kitamura. Empirical likelihood methods with weakly dependent processes. *Ann. Statist.*, 25(5):2084–2102, 1997.

Y. Kitamura and T. Otsu. Bayesian analysis of moment condition models using nonparametric priors. Technical report, Department of Economics: Yale University, 2011.

Y. Kitamura and M. Stutzer. An information-theoretic alternative to generalized method of moments estimation. *Econometrica*, 65(4):861–874, 1997.

Y. K. Kwan. Asymptotic bayesian analysis based on a limited information estimator. *Journal of Econometrics*, 88(1):99–121, 1999.

N. A. Lazar. Bayesian empirical likelihood. *Biometrika*, 90(2):319–326, 2003.

Y. Liao and W. Jiang. Posterior consistency of nonparametric conditional moment restricted models. *The Annals of Statistics*, 39(6):pp. 3003–3031, 2011.

W. K. Newey and R. J. Smith. Higher order properties of gmm and generalized empirical likelihood estimators. *Econometrica*, 72(1):219–255, 2004.

A. B. Owen. Empirical likelihood ratio confidence intervals for a single functional. *Biometrika*, 75(2):237–249, 1988.

A. B. Owen. *Empirical Likelihood*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability, 2001.

J. Qin and J. Lawless. Empirical likelihood and general estimating equations. *The Annals of Statistics*, 22(1):300–325, 03 1994.

J. Rao and C. Wu. Bayesian pseudo-empirical-likelihood intervals for complex surveys. *Journal of the Royal Statistical Society Series B*, 72(4):533–544, 2010.

C. Robert. Méthodes de calcul en analyse bayésienne. In J. F. J.J. Droesbeke and G. Saporta, editors, *Méthodes Bayésiennes en statistique*, pages 149 – 181. TECHNIP, 2002.

S. M. Schennach. Bayesian exponentially tilted empirical likelihood. *Biometrika*, 92(1): 31–46, 2005.

S. M. Schennach. Point estimation with exponentially tilted empirical likelihood. *The Annals of Statistics*, 35(2):634–672, 2007.

M. Shin. Bayesian GMM. Technical report, University of Pennsylvania, 2014.

R. J. Smith. Alternative semi-parametric likelihood approaches to generalised method of moments estimation. *The Economic Journal*, 107(441):503–519, 1997.